

06-07-00

A

**IN THE UNITED STATES PATENT AND TRADEMARK OFFICE**

**CERTIFICATE OF EXPRESS MAILING**

certify that this paper and the documents and/or fees referred to as therein are being deposited with the United States Postal Service on June 05, 2000 in an envelope as "Express Mail Post Office to Addressee" service under 37 CFR §1.10, Mailing Label Number EL560314212US, addressed to Box Patent Application, Assistant Commissioner for Patents, Washington, DC 20231.

Attorney Docket No.: CISC139

First Named Inventor: Martin Cieslak

*Dionna Holmes*  
Dionna Holmes

**UTILITY PATENT APPLICATION TRANSMITTAL (37 CFR. § 1.53(b))**

Assistant Commissioner for Patents  
Box Patent Application  
Washington, DC 20231

☐ Duplicate for  
fee processing

Sir: This is a request for filing a patent application under 37 CFR. § 1.53(b) in the name of inventors:  
**Martin Cieslak, James Aviani, Martin Kagan**

For: **NETWORK CACHE-BASED CONTENT ROUTING**

Application Elements:

- ☒ 21 Pages of Specification, Claims and Abstract  
☒ 4 Sheets of **Informal** Drawings  
☐ 6 Pages Combined Declaration and Power of Attorney (**2 documents**)

Accompanying Application Parts:

- ☒ Assignment and Assignment Recordation Cover Sheet (recording fee of \$40.00 enclosed)  
☐ 37 CFR 3.73(b) Statement by Assignee  
☐ Information Disclosure Statement with Form PTO-1449  
☐ Copies of IDS Citations  
☐ Preliminary Amendment  
☒ Return Receipt Postcard  
☐ Small Entity Statement(s)  
☐ Other:

Fee Calculation (37 CFR § 1.16)

	(Col. 1) NO. FILED	(Col. 2) NO. EXTRA	SMALL ENTITY RATE FEE	OR	LARGE ENTITY RATE FEE
BASIC FEE			\$345 \$	OR	\$690 \$ 690
TOTAL CLAIMS	23 -20 = 3		x 9 = \$	OR	x 18 = \$ 54
INDEP CLAIMS	6 -03 = 3		x 39 = \$	OR	x 78 = \$ 234
[ ] Multiple Dependent Claim Presented			\$130 = \$	OR	\$260 = \$
* If the difference in Col. 1 is less than zero, enter "0" in Col. 2.			Total \$	OR	Total \$ 978

☒ Check No. 2873 in the amount of \$ 1,018.00 is enclosed.

☒ The Commissioner is authorized to charge any fees beyond the amount enclosed which may be required, or to credit any overpayment, to Deposit Account No. 50-0388 (Order No. CISC139).

General Authorization for Petition for Extension of Time (37 CFR §1.136)

☒ Applicants hereby make and generally authorize any Petitions for Extensions of Time as may be needed for any subsequent filings. The Commissioner is also authorized to charge any extension fees under 37 CFR §1.17 as may be needed to Deposit Account No. 50-0388 (Order No. CISC139).

☒ Please send correspondence to the following address:

**Customer Number 022434**  
**BEYER WEAVER & THOMAS, LLP**  
P.O. Box 130  
Mountain View, CA 94042-0130  
Telephone (510) 843-6200  
Fax (510) 843-6203



Date: June 5, 2000

Joseph M. Villeneuve  
**Joseph M. Villeneuve**  
Registration No. **37,460**

PATENT APPLICATION

NETWORK CACHE-BASED CONTENT ROUTING

Inventor(s):           Martin Cieslak of  
                              Fremont, California  
                              United States citizen

                              James Aviani of  
                              Santa Barbara, California  
                              United States citizen

                              Martin Kagan of  
                              Burlingame, California  
                              United States citizen

Assignee:               Cisco Technology Inc.

BEYER WEAVER & THOMAS, LLP  
P.O. Box #130  
Mountain View, CA 94042-0130  
(510) 843-6200

005090" 22088560

## NETWORK CACHE-BASED CONTENT ROUTING

### BACKGROUND OF THE INVENTION

The present invention relates to transmission of data in a network environment. More specifically, the present invention relates to methods and apparatus for improving the efficiency with which data are transmitted over the Internet. Still more specifically, the present invention provides techniques by which routing decisions for particular data packets are made with reference to the nature of the data being transmitted, e.g., content routing.

Generally speaking, when a client platform communicates with some remote server, whether via the Internet or an intranet, it crafts a data packet which defines a TCP connection between the two hosts, i.e., the client platform and the destination server. More specifically, the data packet has headers which include the destination IP address, the destination port, the source IP address, the source port, and the protocol type. The destination IP address might be the address of a well known World Wide Web (WWW) search engine such as, for example, Yahoo, in which case, the protocol would be TCP and the destination port would be port 80, a well known port for http and the WWW. The source IP address would, of course, be the IP address for the client platform and the source port would be one of the TCP ports selected by the client. These five pieces of information define the TCP connection.

Given the increase of traffic on the World Wide Web and the growing bandwidth demands of ever more sophisticated multimedia content, there has been constant pressure to find more efficient ways to service data requests than opening direct TCP connections between a requesting client and the primary repository for the desired data. Interestingly,

one technique for increasing the efficiency with which data requests are serviced came about as the result of the development of network firewalls in response to security concerns. In the early development of such security measures, proxy servers were employed as firewalls to protect networks and their client machines from corruption by undesirable content and unauthorized access from the outside world. Proxy servers were originally based on Unix machines because that was the prevalent technology at the time. This model was generalized with the advent of SOCKS which was essentially a daemon on a Unix machine. Software on a client platform on the network protected by the firewall was specially configured to communicate with the resident demon which then made the connection to a destination platform at the client's request. The demon then passed information back and forth between the client and destination platforms acting as an intermediary or "proxy."

Not only did this model provide the desired protection for the client's network, it gave the entire network the IP address of the proxy server, therefore simplifying the problem of addressing of data packets to an increasing number of users. Moreover, because of the storage capability of the proxy server, information retrieved from remote servers could be stored rather than simply passed through to the requesting platform. This storage capability was quickly recognized as a means by which access to the World Wide Web could be accelerated. That is, by storing frequently requested data, subsequent requests for the same data could be serviced without having to retrieve the requested data from its original remote source. Currently, most Internet service providers (ISPs) accelerate access to their web sites using proxy servers.

Unfortunately, interaction with such proxy servers is not transparent, requiring each end user to select the appropriate proxy configuration in his or her browser to allow the browser to communicate with the proxy server. For the large ISPs with millions of customers there is significant overhead associated with handling tech support calls from

customers who have no idea what a proxy configuration is. Additional overhead is associated with the fact that different proxy configurations must be provided for different customer operating systems. The considerable economic expense represented by this overhead offsets the benefits derived from providing accelerated access to the World Wide Web. Another problem arises as the number of WWW users increases. That is, as the number of customers for each ISP increases, the number of proxy servers required to service the growing customer base also increases. This, in turn, presents the problem of allocating packet traffic among multiple proxy servers.

Another technique for increasing the efficiency with which data requests are serviced is described in commonly assigned, copending U.S. Patent Application No. 08/946,867 for METHOD AND APPARATUS FOR FACILITATING NETWORK DATA TRANSMISSIONS filed October 8, 1997, the entirety of which is incorporated herein by reference for all purposes. The invention described in that copending application represents an improvement over the proxy server model which is transparent to end users, high performance, and fault tolerant. By altering the operating system code of an existing router, the router is enabled to redirect data traffic of a particular protocol intended for a specified port, e.g., TCP with port 80, to one or more caching engines connected to the router via an interface having sufficient bandwidth such as, for example, a 100baseT interface. If there are multiple caching engines connected to the cache-enabled router, the router selects from among the available caching engines for a particular request based on a simple algorithm according to which a particular group of addresses is associated with each caching engine.

The caching engine to which the request is re-routed "spoofs" the requested destination platform and accepts the request on its behalf via a standard TCP connection established by the cache-enable router. If the requested information is already stored in the caching engine, i.e., a cache "hit" occurs, it is transmitted to the requesting platform with a

header indicating its source as the destination platform. If the requested information is not in the caching engine, i.e., a cache "miss" occurs, the caching engine opens a direct TCP connection with the destination platform, downloads the information, stores it for future use, and transmits it to the requesting platform. All of this is transparent to the user at the requesting platform which operates exactly as if it were communicating with the destination platform. Thus, the need for configuring the requesting platform to suit a particular proxy configuration is eliminated along with the associated overhead. Moreover, traffic may be easily allocated among as many caching engines as become necessary. Thus, content caching provides a way to compensate for the bandwidth limitations discussed above.

Another specific embodiment of a packet redirection protocol which may be used to implement such a network caching technique is described in copending, commonly assigned, U.S. Provisional Patent Application No. 60/168,862 for METHOD AND APPARATUS FOR REDIRECTING NETWORK TRAFFIC filed December 2, 1999, the entirety of which is incorporated herein by reference for all purposes. According to a specific embodiment described in that application, the network caches have the capability of determining that particular redirected packets should be transmitted back to the redirecting router and reinserted into the original traffic flow. This may be done in a manner transparent to the source or destination of the packets. An example of a case in which packets would need to be reinserted in the original flow might be where the cache recognizes the source and destination pairs identified by the packets as corresponding to a connection requiring IP-based authentication. Another example would be where the cache is overloaded and is currently unable to handle all of the redirected traffic.

A technique conceptually related to content caching referred to as content co-location has also been developed to provide faster access to certain types of data. That is, a large amount of Internet content resides on machines in the U.S. Client machines in remote

geographic locations such as New Zealand or Australia typically have two ways of retrieving such content, i.e., via terrestrial links (bidirectional and fast, but expensive) and satellite links (mostly unidirectional and relatively cheap, but slow).

Content co-location provides one way to facilitate faster access to such content for geographically remote users. According to this technique, servers are provided in the geographically remote areas which store at least some of the content of the original site, e.g., image files. Akamai and Sandpiper provide such “co-loc” facilities to content providers. When a client requests a page containing such objects, the html refers to the objects by object names which correspond to the co-loc provider. These object names are then resolved by the client’s browser to an IP address corresponding to the closest co-loc facility.

While content co-location techniques have proven to be relatively effective, content providers pay a huge premium for this service. In addition, implementation of content co-location requires some relatively complex and computationally intense protocols. For example, coherency must be maintained between content at co-loc facilities and the content at the original site. In addition, browsers must constantly perform the object name resolution when parsing ordinary html to access the correct co-location facility.

Another limitation associated with various network devices, e.g., routers, switches, gateways, and some caches, is that routing decisions are made only with reference to the lower level network layers, i.e., the physical, data link, and network layers. Fig. 1 illustrates the relationship of one such network device, i.e., a gateway, to the various network layers. Assume that the user application layer protocol 102A in host A sends an application protocol data unit (PDU) to an application layer protocol 102B in host B, such as, for example, a file transfer system. The file transfer software performs a variety of functions and sends file records to the user. In many systems, the operations at host B are known as *server* operations and the operations at host A are known as *client* operations.



As indicated by the downward arrows in the protocol stack at host A, this unit is passed to the transport layer protocol 104A, which performs a variety of operations and adds a header to the PDU passed to it. At this point, the unit of data is often referred to as a *segment*. The PDU from the upper layers is considered to be data to the transport layer.

5       Next, the transport layer passes the segment to the network layer 106A, also called the IP *layer*, which again performs specific services and appends a header. This unit (now called a *datagram* in internet terms) is passed down to the lower layers. Here, the data link layer adds its header as well as a trailer, and the data unit (now called a *frame*) is launched into subnetwork 110 by the physical layer 108A. Of course, if host B sends data to host A,  
10      the process is reversed and the direction of the arrows is changed.

Internet protocols are typically unaware of what goes on inside the network. The network manager is free to manipulate and manage the PDU in any manner necessary. In some instances, however, the internet PDU (data and headers) remains unchanged as it is transmitted through the subnet. In Fig. 1, it emerges at the gateway where it is processed  
15      through the lower layers 114 and passed to the IP (network) layer 112. Here, routing decisions are made based only on the destination address provided by the host computer. After these routing decisions have been made, the PDU is passed to the communications link connected to the appropriate subnetwork (comprising the lower layers). The PDU is re-encapsulated into the data link layer frame and passed to the next subnetwork 116, where it  
20      finally arrives at the destination host.

One of the limitations inherent in the traffic routing approach described above is that the routing decisions are made with reference only to information in the lower network layers. Because data transmissions in these layers are virtually indistinguishable from one another it is difficult if not impossible to make intelligent routing decisions based on the type

of data being transmitted or requested, or any other high level information. As a result, opportunities for improving the efficiency with which networks operate are lost.

In view of the foregoing, there is a need for techniques by which network devices can make intelligent routing decisions about specific types of data traffic which are less

5 expensive in terms of network resources than currently available solutions.

005090" 22037560  
09120027 060500

## SUMMARY OF THE INVENTION

According to the present invention, network devices are enabled to make intelligent routing decisions based upon any of a number of parameters or characteristics associated with specific data traffic. According to a specific embodiment of the invention, the  
5 parameters or characteristics of interest are associated with the application layer. According to another specific embodiment, the parameters or characteristics of interest relate to information which is specific to a network cache making the routing decisions.

Thus, the present invention provides a computer-implemented method for routing data traffic in a network having a plurality of layers including physical, data link, and  
10 network layers. The data traffic is received and one of a plurality of routing options for the data traffic is selected with reference to information outside of the physical, data link, and network layers. The data traffic is then routed according to the selected routing option.

According to one embodiment, the data traffic is received and one of a plurality of routing options for the data traffic is selected with reference to information associated with  
15 the network's application layer. In another embodiment where the data traffic has been redirected to a network cache, the data traffic are received by the network cache which selects one of a plurality of routing options for the data traffic with reference to information about the data traffic accessible by the network cache.

A further understanding of the nature and advantages of the present invention may be  
20 realized by reference to the remaining portions of the specification and the drawings.

## BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1 is a block diagram illustrating conventional data transmissions in a wide area network;

Fig. 2 is a block diagram of a network environment in which specific embodiments  
5 of the present invention may be implemented;

Fig. 3 is a block diagram of a network cache for use with specific embodiments of the present invention; and

Fig. 4 is a flowchart illustrating the caching of objects according to a specific embodiment of the present invention.

## DETAILED DESCRIPTION OF SPECIFIC EMBODIMENTS

Fig. 2 is a block diagram of a network environment in which specific embodiments of the present invention may be implemented. According to a first specific embodiment, a request for content in the United States is received by a network cache 210 in Australia.

5 Based on information associated with the application layer, cache 210 decides whether to use a terrestrial link or a satellite link to route the request. The selected link will thus determine the routing path for the response to the request. According to one embodiment, the determination is made by parsing the request in the application layer and then selecting a socket option for the outgoing request accordingly. In the figure, the terrestrial link is shown  
 10 comprising router 212, a wide area network 214 of arbitrary size, and router 216. The satellite link is shown comprising router 218, wide area network 220, satellite link 222, satellite 224, satellite link 226, wide area network 228, and router 230. Each of the links is shown connecting with a cache 232 although this is not necessary for implementation of this embodiment. That is, the links may connect directly with the destination machine or an  
 15 associated router.

According to a more specific embodiment, the decision between the terrestrial and satellite links is based on information which correlates to the relative sizes of the objects being requested. That is, html objects are typically fairly small as compared to graphically intense embedded objects. In addition, html objects typically anchor or form the foundation  
 20 for web pages. When a user is downloading such a page, they will tolerate a slower download of the graphics, but they need the page's foundation, i.e., the html, right away. Therefore, according to this embodiment, the html is requested via the fast terrestrial link and the embedded graphics via the slower satellite link. According to a still more specific embodiment, the type of object is determined (i.e., the routing decision is made) with  
 25 reference to the suffix of the URL (e.g., \*.htm vs. \*.gif).

The embodiments described above use the satellite uplink to request the larger objects. That is, because the decision making is performed on the side of the links which is remote from the requested content, satellite link 222 must be used to both transmit the request and receive the requested data. Because providing this uplink capability at satellite link 222 can be expensive, a second embodiment is provided in which cache 232 on the content side of the links makes the decisions. That is, when cache 232 receives data traffic, it routes the traffic using, for example, the criteria described above with reference to cache 210. The difference is that, because cache 232 is on the content side of the links, no satellite uplink needs to be provided. Thus, for example, as with the embodiment described above, because of the application layer characteristics of data traffic received by cache 232, large objects (e.g., \*.gif objects) may be sent via the satellite link while small objects (\*.htm objects) may be sent via the terrestrial link.

It will be understood that the network configurations and links shown in Fig. 2 are merely used for illustration purposes and that a wide variety of configurations may be employed without departing from the scope of the invention. In addition, the routing decisions enabled by the present invention are not limited to a choice between a terrestrial link and a satellite link, or based only on application layer information as described above. Rather, the present invention enables intelligent selection from among a plurality of routing options based on any combination of information associated with the application layer and any of the other network layers. Put another way, based upon the request, the cheapest of a plurality of links is selected for transmitting the request and thus the object. The determination of what is cheapest is a multi-parameter determination which looks at, for example, speed, economic cost, specific user's need/level of access, remote ISP use agreements (e.g., many people overseas have multiple ISPs for which the cost of use increases dramatically at certain levels thereby making it reasonable to switch over to

another ISP), etc. Effectively, the present invention enables so called content routing, i.e., the notion that objects are routed based on the type of content requested.

The information upon which the routing determination is based may vary widely. That is, network caches designed according to the present invention may incorporate routing rules based on any information associated with the higher network layers. For example, a request or a requested object may be routed according to whether the requested content is cacheable or uncacheable, ascii or binary, HTTP or ICP, regular requests or forced reloads, static web page requests or browser-based applications. For example, an object may be determined to be cacheable by looking at the HTTP headers or URL suffixes. In addition, whether an object is ascii or binary may be determined with reference to specific file suffixes, e.g., .htm vs. .gif or .jpg. An object can be differentiated as HTTP or ICP with reference to the transmission protocol used (i.e., TCP vs. UDP) or the port. A regular request may be distinguished from a forced reload with reference to the appropriate HTTP header.

Fig. 3 is a block diagram of a network cache such as, for example, cache 210 or cache 232 of Fig. 2. A central processor 302 controls operation of cache 210/232 and its various subsystems using system memory 304 and bus 306. Data objects are stored in cache memory 308 which, in a specific embodiment, comprises multiple queues of volatile RAM 310. According to various embodiments, memory 308 may comprise one or more nonvolatile disk drives. According to yet other embodiments, memory 308 may comprise any combination of volatile and nonvolatile memory. According to a specific embodiment, a nonvolatile disk drive 311 is provided as additional storage for cached objects.

A network interface 312 enables communication with external devices. Portions of memory 310 may also be employed for other purposes such as, for example, storing software code for directing the operation of various functionalities of cache 210/232. Alternatively,

program instructions for execution by processor 302 directing operation of the functionalities of cache 210/232 may be stored in a separate program memory 305. It will be understood that the cache architecture shown in Fig. 3 is merely illustrative and should not be construed to limit the scope of the present invention. That is, any of a wide variety of cache architectures may be employed to implement the present invention.

Network caches are in a unique position to have certain information about specific data traffic which is unavailable to other network devices. Some of this information is useful with regard to making intelligent routing decisions for the associated traffic.

According to a specific embodiment of the present invention, when a network cache receives a redirected request for an object from a cache-enabled router, the cache identifies the redirecting router from the encapsulation header and, if the cache needs to issue its own request for that object, it can route the new request through the redirecting router to avoid disruption of an independent load balancing scheme. This is important in caching systems in which multiple routers are logically connected to multiple caches.

Fig. 4 is a flowchart illustrating the routing of data traffic according to a specific embodiment of the present invention. When a packet is redirected to a cache designed according to the present invention (402), information associated with the packet is parsed (404) to determine which of a plurality of routing options are to be selected (406). The packet is then routed according to the selected routing option (408). According to a specific embodiment and as described above, the information associated with the packet is application layer information. According to yet other embodiments, cache specific information associated with particular data traffic is parsed for the routing determination.

While the invention has been particularly shown and described with reference to specific embodiments thereof, it will be understood by those skilled in the art that changes in the form and details of the disclosed embodiments may be made without departing from the



[illegible]

WHAT IS CLAIMED IS:

1. A computer-implemented method for routing data traffic in a network having a plurality of network layers including an application layer, the method comprising:

5 receiving the data traffic;

selecting one of a plurality of routing options for the data traffic with reference to information associated with the application layer; and

routing the data traffic according to the selected routing option.

10 2. The method of claim 1 wherein the data traffic has been redirected from an original destination according to a caching protocol.

3. The method of claim 1 wherein the data traffic comprises a request from a source platform to a destination platform.

15 4. The method of claim 1 wherein the data traffic comprises a response to a request, the request being from a source platform to a destination platform.

5. The method of claim 1 further comprising parsing the information associated with the application layer.

20 6. The method of claim 5 wherein the information comprises a URL associated with the data traffic.

7. The method of claim 6 wherein the information comprises a suffix associated with the URL.

8. The method of claim 7 wherein parsing the information comprises  
5 determining whether the suffix associated with the URL indicates one of a plurality of MIME types.

9. The method of claim 8 wherein the plurality of MIME types comprises \*.gif, \*.jpg, \*.pdf, \*.mpX, and \*.htm.

10. The method of claim 5 wherein parsing the information comprises  
10 determining whether the data traffic relates to ascii or binary data objects.

11. The method of claim 1 wherein selecting one of the plurality of options  
15 comprises setting one of a plurality of socket options for the data traffic.

12. The method of claim 11 wherein the plurality of socket options include a first link and a second link, the first link socket option being selected for a first type of data traffic and the second link socket option being selected for a second type of data traffic.

13. The method of claim 12 wherein the first and second links comprise land and satellite links, respectively.

14. The method of claim 12 wherein the first and second types of data comprise  
25 ascii and binary data, respectively.

15. A computer program product comprising a computer readable medium having computer program instructions stored therein for implementing the method of claim 1.

5 16. A computer-implemented method for routing data traffic in a network which has been redirected to a network cache, the method comprising:

receiving the data traffic with the network cache;

selecting one of a plurality of routing options for the data traffic with reference to information about the data traffic accessible by the network cache; and

10 routing the data traffic according to the selected routing option.

17. The method of claim 16 wherein the information relates to whether a data object associated with the data traffic is cacheable.

15 18. The method of claim 16 wherein the information relates to whether the data traffic comprises a forced reload.

19. A computer program product comprising a computer readable medium having computer program instructions stored therein for implementing the method of claim 16.

20 20. A computer-implemented method for routing data traffic in a network having a plurality of layers including physical, data link, and network layers, the method comprising:

receiving the data traffic;

selecting one of a plurality of routing options for the data traffic with reference to information outside of the physical, data link, and network layers; and routing the data traffic according to the selected routing option.

5           21.     A network cache for operating in a network having a plurality of layers including an application layer, comprising:

cache memory for storing a plurality of objects; and

an operating system which is operable to:

receive redirected data traffic;

10           select one of a plurality of routing options for the data traffic with reference to information associated with the application layer; and

route the data traffic according to the selected routing option.

22.     A network cache, comprising:

15           cache memory for storing a plurality of objects; and

an operating system which is operable to:

receiving redirected data traffic;

select one of a plurality of routing options for the data traffic with reference to information about the data traffic accessible by the network cache; and

20           route the data traffic according to the selected routing option.

23.     A network cache for operating in a network having a plurality of layers including physical, data link, and network layers, comprising:

cache memory for storing a plurality of objects; and

25           an operating system which is operable to:

receive redirected data traffic;

select one of a plurality of routing options for the data traffic with reference to information outside of the physical, data link, and network layers; and

route the data traffic according to the selected routing option.

## NETWORK CACHE-BASED CONTENT ROUTING

## ABSTRACT OF THE DISCLOSURE

Methods and apparatus are described for routing data traffic in a network having a plurality of layers including physical, data link, and network layers. The data traffic is received and one of a plurality of routing options for the data traffic is selected with reference to information outside of the physical, data link, and network layers. The data traffic is then routed according to the selected routing option.

005090 22083560

10




Fig. 1






Fig. 2

210/232

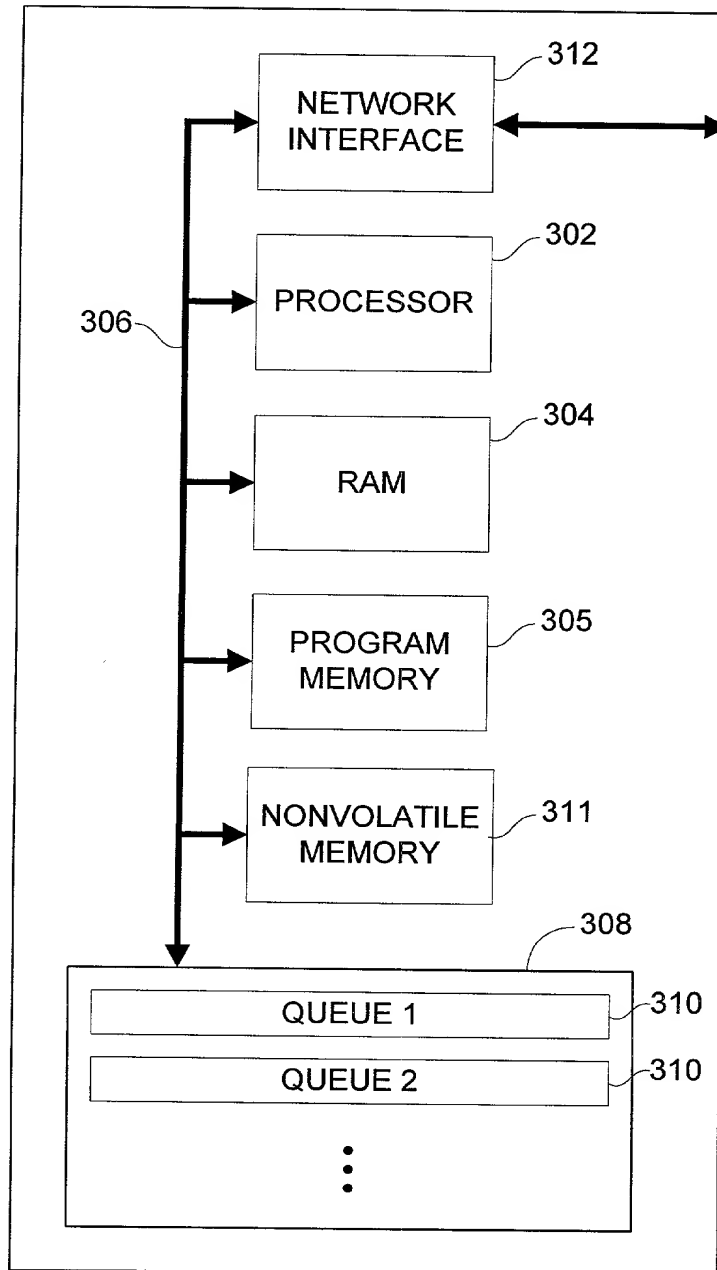


Fig. 3

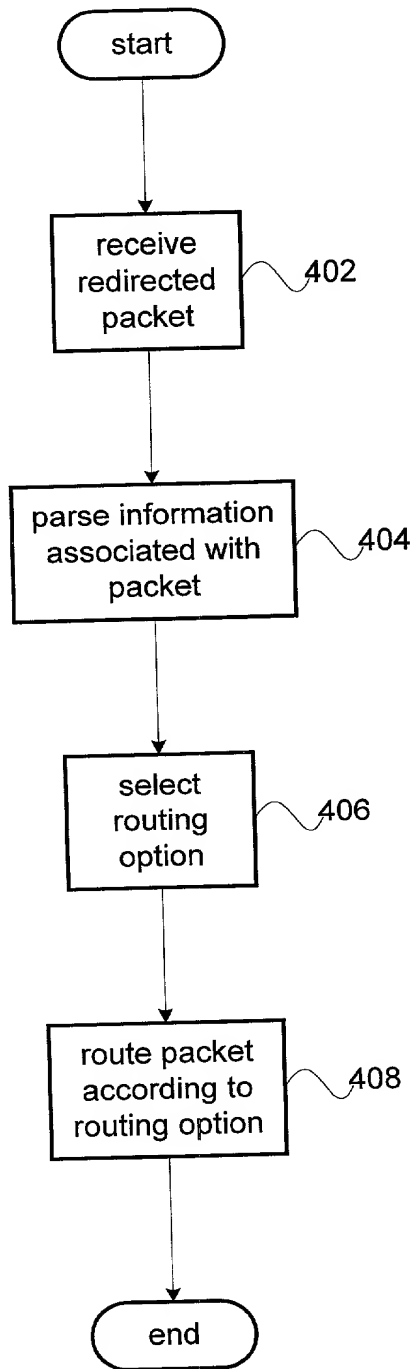


Fig. 4

# DECLARATION AND POWER OF ATTORNEY FOR ORIGINAL U.S. PATENT APPLICATION

Attorney's Docket No. **CISCP139/1594**

As a below-named inventor, I/we hereby declare that:

My residence, post office address and citizenship are as stated below next to my name.

I believe that I am the original, first and sole inventor (if only one name is listed below) or an original, first and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled: **NETWORK CACHE-BASED CONTENT ROUTING**, the specification of which,

(check one)

1. ☒ is attached hereto.

2. ☐ was filed on \_\_\_\_\_ as  
U.S. Application No. \_\_\_\_\_  
and was amended on \_\_\_\_\_.

3. ☐ was filed on \_\_\_\_\_ as  
International PCT Application No. \_\_\_\_\_  
and was amended on \_\_\_\_\_.

hereby state that I have reviewed and understand the contents of the above-identified specification, including the claims, as amended by any amendment referred to above.

acknowledge the duty to disclose information which is material to the examination of this application in accordance with Title 37, CFR § 1.56.

hereby claim foreign priority benefits under Title 35, United States code, § 119(a)-(d) or § 365(b) of any foreign application(s) for patent or inventor's certificate, or § 365(a) of any PCT International application which designated at least one country other than the United States, listed below and have identified below, by checking the box, any foreign application for patent or inventor's certificate, or PCT International application having a filing date before that of the application on which priority is claimed:

## Prior Foreign Application(s)

Priority Benefits Claimed?

☐ Yes ☐ No

(Appl. No.)

(Country)

(Filing Date)

I hereby claim the benefit under 35 U.S.C. §119(e) of any United States provisional application(s) listed below:

## Prior Provisional Application(s)

(Application No.)

(Filing Date)

I hereby claim the benefit under Title 35, United States Code, § 120 of any United States application(s), or § 365(c) of any PCT International application designating the United States, listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States or PCT International application in the manner provided by the first paragraph of Title 35, United States Code, § 112, I acknowledge the duty to disclose information which is material to patentability as defined in Title 37, Code of Federal Regulations, § 1.56 which became available between the filing date of the prior application and the national or PCT international filing date of this application:

**Prior U.S. Application(s)**

(Application No.) (Filing Date) (Status - patented, pending, abandoned)

And I hereby appoint the law firm of **Beyer & Weaver, LLP** and all practitioners who are associated with the Customer Number 022434 as my principal attorneys to prosecute this application and to transact all business in the Patent and Trademark Office connected therewith.

Direct Correspondence To:

**Customer Number: 022434**  
BEYER & WEAVER, LLP  
P.O. Box 61059  
Palo Alto, CA 94306



Direct Telephone Calls To:

Joseph M. Villeneuve at telephone number (510) 843-6200

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under section 1001 of Title 18 of the United States Code, and that such willful false statements may jeopardize the validity of the application or any patent issuing thereon.

Typewritten Full Name of

First Joint Inventor:

Martin Cieslak

Citizenship: US

Inventor's signature:

Date of Signature: 3/31/00

Residence (City):

Fremont

(State/Country) CA/USA

Post Office Address:

1074 Canyon Creek Terrace, Fremont, CA 94536

Typewritten Full Name of

Second Joint Inventor:

James Avlani

Citizenship: US

Inventor's signature:

Date of Signature: 4/27/2000

Residence (City):

Santa Barbara

(State/Country) CA/USA

Post Office Address:

170 Olive Mill Road, Santa Barbara, CA 93108

Typewritten Full Name of  
Third Joint Inventor:

Martin Kagan

Citizenship: US

**Inventor's signature:**

\_\_\_\_\_

**Date of Signature:** \_\_\_\_\_

Residence (City):

Burlingame

(State/Country) CA/USA

Post Office Address:

1400 Floribunda Ave., Apt. 201 , Burlingame, CA 94010

005050" 20080560

# DECLARATION AND POWER OF ATTORNEY FOR ORIGINAL U.S. PATENT APPLICATION

Attorney's Docket No. **CISCP139/1594**

As a below-named inventor, I/we hereby declare that:

My residence, post office address and citizenship are as stated below next to my name.

I believe that I am the original, first and sole inventor (if only one name is listed below) or an original, first and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled: **NETWORK CACHE-BASED CONTENT ROUTING**, the specification of which,

(check one)

1. ☒ is attached hereto.

2. ☐ was filed on \_\_\_\_\_ as  
U.S. Application No. \_\_\_\_\_  
and was amended on \_\_\_\_\_.

3. ☐ was filed on \_\_\_\_\_ as  
International PCT Application No. \_\_\_\_\_  
and was amended on \_\_\_\_\_.

I hereby state that I have reviewed and understand the contents of the above-identified specification, including the claims, as amended by any amendment referred to above.

I acknowledge the duty to disclose information which is material to the examination of this application in accordance with Title 37, CFR § 1.56.

I hereby claim foreign priority benefits under Title 35, United States code, § 119(a)-(d) or § 365(b) of any foreign application(s) for patent or inventor's certificate, or § 365(a) of any PCT International application which designated at least one country other than the United States, listed below and have identified below, by checking the box, any foreign application for patent or inventor's certificate, or PCT International application having a filing date before that of the application on which priority is claimed:

## Prior Foreign Application(s)

Priority Benefits Claimed?

☐ Yes ☐ No

\_\_\_\_\_  
(Appl. No.)

\_\_\_\_\_  
(Country)

\_\_\_\_\_  
(Filing Date)

I hereby claim the benefit under 35 U.S.C. §119(e) of any United States provisional application(s) listed below:

## Prior Provisional Application(s)

\_\_\_\_\_  
(Application No.)

\_\_\_\_\_  
(Filing Date)

I hereby claim the benefit under Title 35, United States Code, § 120 of any United States application(s), or § 365(c) of any PCT International application designating the United States, listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States or PCT International application in the manner provided by the first paragraph of Title 35, United States Code, § 112, I acknowledge the duty to disclose information which is material to patentability as defined in Title 37, Code of Federal Regulations, § 1.56 which became available between the filing date of the prior application and the national or PCT international filing date of this application:

**Prior U.S. Application(s)**

<u>                    </u> (Application No.)	<u>                    </u> (Filing Date)	<u>                    </u> (Status - patented, pending, abandoned)
--	--	--

And I hereby appoint the law firm of **Beyer & Weaver, LLP** and all practitioners who are associated with the Customer Number 022434 as my principal attorneys to prosecute this application and to transact all business in the Patent and Trademark Office connected therewith.

Direct Correspondence To:

**Customer Number: 022434**  
**BEYER & WEAVER, LLP**  
P.O. Box 61059  
Palo Alto, CA 94306



Direct Telephone Calls To:

**Joseph M. Villeneuve at telephone number (510) 843-6200**

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under section 1001 of Title 18 of the United States Code, and that such willful false statements may jeopardize the validity of the application or any patent issuing thereon.

Typewritten Full Name of  
First Joint Inventor:

Martin Cieslak

Citizenship: US

**Inventor's signature:**

**Date of Signature:**                     

Residence (City):

Fremont

(State/Country) CA/USA

Post Office Address:

1074 Canyon Creek Terrace, Fremont, CA 94536

Typewritten Full Name of  
Second Joint Inventor:

James Aviani

Citizenship: US

**Inventor's signature:**

**Date of Signature:**                     

Residence (City):

Santa Barbara

(State/Country) CA/USA

Post Office Address:

170 Olive Mill Road, Santa Barbara, CA 93108




Typewritten Full Name of  
Third Joint Inventor:

Martin Kagan

Citizenship: US

**Inventor's signature:**



**Date of Signature:** 4/26/00

Residence (City):

Burlingame

(State/Country) CA/USA

Post Office Address:

1400 Floribunda Ave., Apt. 201 , Burlingame, CA 94010

005090 20000560